Document made available under the Patent Cooperation Treaty (PCT)

International application number: PCT/US2003/017608

International filing date: 04 June 2003 (04.06.2003)

Document type: Certified copy of priority document

Document details: Country/Office: US
Number: 60/392,779
Filing date: 01 July 2002 (01.07.2002)

Date of receipt at the International Bureau: 25 April 2007 (25.04.2007)

Remark: Priority document submitted or transmitted to the International Bureau in compliance with Rule 17.1(a) or (b)

1600539

THE UNITED STATES OF AMERICA

TO ALL TO WHOM THESE PRESENTS SHALL COME:

UNITED STATES DEPARTMENT OF COMMERCE

United States Patent and Trademark Office

April 19, 2007

THIS IS TO CERTIFY THAT ANNEXED HERETO IS A TRUE COPY FROM THE RECORDS OF THE UNITED STATES PATENT AND TRADEMARK OFFICE OF THOSE PAPERS OF THE BELOW IDENTIFIED PATENT APPLICATION THAT MET THE REQUIREMENTS TO BE GRANTED A FILING DATE.

APPLICATION NUMBER: *60/392,779*
FILING DATE: *July 01, 2002*
RELATED PCT APPLICATION NUMBER: *PCT/US03/17608*

THE COUNTRY CODE AND NUMBER OF YOUR PRIORITY APPLICATION, TO BE USED FOR FILING ABROAD UNDER THE PARIS CONVENTION, IS *US60/392,779*

Certified by

Jon W. Dudas

Under Secretary of Commerce
for Intellectual Property
and Director of the United States
Patent and Trademark Office

07-02-02 60392779 .020102 AJRW

Exp. Mail Label: EV 104020697 US

Please type a plus sign (+) inside this box → +

PTO/SB/16 (8-00)
Approved for use through 10/31/2002 OMB 0651-0032
Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

PROVISIONAL APPLICATION FOR PATENT COVER SHEET

This is a request for filing a PROVISIONAL APPLICATION FOR PATENT under 37 CFR 1.53 (c).

INVENTOR(S)		
Given Name (first and middle [if any])	Family Name or Surname	Residence (City and either State or Foreign Country)
David Aaron	Crowther	Hillsboro, Oregon 97123
Ravi	Reddy	Beaverton, Oregon 97006
Andy Eugene	Adkins	Portland, Oregon 97225

☒ *Additional inventors are being named on the* _____ *separately numbered sheets attached hereto*

TITLE OF THE INVENTION (280 characters max)

Heterogeneous Low Level Disk Storage Management Tool

CORRESPONDENCE ADDRESS

Direct all correspondence to:

☐ Customer Number [ ] → *Place Customer Number Bar Code Label here*

OR *Type Customer Number here*

☒ Firm *or* Individual Name	THOMSON MULTIMEDIA LICENSING INC.				
Address	PATENT OPERATIONS				
Address	P. O. BOX 5312				
City	PRINCETON	State	NJ	ZIP	08543-5312
Country	USA	Telephone	609-734-	Fax	609-734-9700

ENCLOSED APPLICATION PARTS *(check all that apply)*

☒ Specification *Number of Pages* 2 ☐ CD(s), Number [ ]

☒ Drawing(s) *Number of Sheets* 1 ☒ Other (specify) [ ] Addendum A & B Attached

☐ Application Data Sheet. See 37 CFR 1.76

METHOD OF PAYMENT OF FILING FEES FOR THIS PROVISIONAL APPLICATION FOR PATENT *(check one)*

☐ Applicant claims small entity status. See 37 CFR 1.27.

☐ A check or money order is enclosed to cover the filing fees

FILING FEE AMOUNT ($)

☒ The Commissioner is hereby authorized to charge filing fees or credit any overpayment to Deposit Account Number: 07-0832 — 160.00

☐ Payment by credit card. Form PTO-2038 is attached.

The invention was made by an agency of the United States Government or under a contract with an agency of the United States Government.

☒ No.

☐ Yes, the name of the U.S. Government agency and the Government contract number are: _____.

Respectfully submitted,

SIGNATURE _____ Date 7-1-02

TYPED or PRINTED NAME Joel M. Fogelson REGISTRATION NO. *(if appropriate)* 43,613

TELEPHONE 609-734-9534 Docket Number: PU020329

USE ONLY FOR FILING A PROVISIONAL APPLICATION FOR PATENT

This collection of information is required by 37 CFR 1.51. The information is used by the public to file (and by the PTO to process) a provisional application Confidentiality is governed by 35 U.S.C 122 and 37 CFR 1 14 This collection is estimated to take 8 hours to complete, including gathering, preparing, and submitting the complete provisional application to the PTO Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U S Patent and Trademark Office, U.S. Department of Commerce, Washington, D.C., 20231. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Box Provisional Application, Assistant Commissioner for Patents, Washington, D.C. 20231.

Exp. Mail Label: EV 104020697US

PROVISIONAL APPLICATION COVER SHEET

Additional Page

PTO/SB/16 (8-00)
Approved for use through 10/31/2002. OMB 0651-0032
Patent and Trademark Office, U S DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number

Docket Number	PU020329	Type a plus sign (+) inside this box →	+

INVENTOR(S)/APPLICANT(S)

Given Name (first and middle [if any])	Family or Surname	Residence (City and either State or Foreign Country)
Nanyu	Cao	Portland, Oregon 97229

Number 2 of 2

WARNING: Information on this form may become public. Credit card information should not be included on this form. Provide credit card information and authorization on PTO-2038.

PU020329

Exp. Mail Label: EV 104020697 US

Heterogeneous Low Level Disk Storage Management Tool

This invention consists of a storage management tool with a unique GUI that abstracts the lower level configuration complexities of disparate disk subsystems, all on the same Media Area Network (MAN)

The problem is that of configuring, *at a low level*, multiple heterogeneous disk subsystems in a uniform manner using an interface as a GUI (FIG. 1). Each disk subsystem (RAID) vendor produces unique devices that require specific configuration commands. Some of these include:

- binding a group of physical disk drives into a Logical Unit Number (LUN)
- unbinding LUNs
- designating physical disk drives as Hot Spares
- setting system clock
- loading vendor specific firmware to individual controllers and/or drives
- loading vendor specific controller/disk settings
- blinking disk drive LEDs for identification purposes
- identifying state information
- setting Fibre Channel IDs for each device

When originally presented with the problem of managing *at a low level* multiple heterogeneous disk subsystems on the same network, it was decided that we would look at what was out there already. One can discover that there are a number of software applications that manage heterogeneous networks *at a high level*, but nothing that would allow configuration. at the level we needed across multiple disk subsystem vendors. One can also discover that each disk subsystem vendor provides a discreet management tool that works for their system only, but no tools were available to manage multiple heterogeneous disk subsystems on the same network. It was decided that we would implement our own tool with a uniform GUI that masks the complexities of each individual disk subsystem vendor, and support multiple of these using the same unique and simplified approach

The invention consists of a heterogeneous *low level* disk storage management tool that allows the configuration *at a low level* multiple of desparate disk subsystems on the same network. The GUI consists of a tree control with vendor specific configuration options appearing when you right click on each vendor controller, LUN, or disk.

The graphic above depicts the main component of the GUI. Various other popup boxes and informational boxes appear automatically when selecting a command (see Appendix B). The tool is simple to use, notwithstanding the underlying complexities of each disk subsystem. The current embodiment includes support for three disparate RAID disk subsystems (EMC/Clariion, Ciprico, and NEC). However, future embodiments could include support for countless different disk subsystem vendors. This tool may be used on discrete systems, or heterogeneous Storage Area Networks. Currently it will run on any standard Windows operating system with a standard FC/SCSI

adapter. Future embodiments of this tool could run on multiple operating system platforms, as Linux, Lindows, MacOS. Much of the novelty of this utility lies in the inherently intuitive GUI. Additionally, the invention may be adapted to a local media area network for use in the home or office, where the drive systems are contained within different media devices (for a home system) as a personal video player, home computer, DVD-RW, and the like.

Appendix A

Software Specification Document. This document refers to the conception of the invention.

Appendix B

This document gives detailed usage and operational information, but only as it applies to a discrete system.

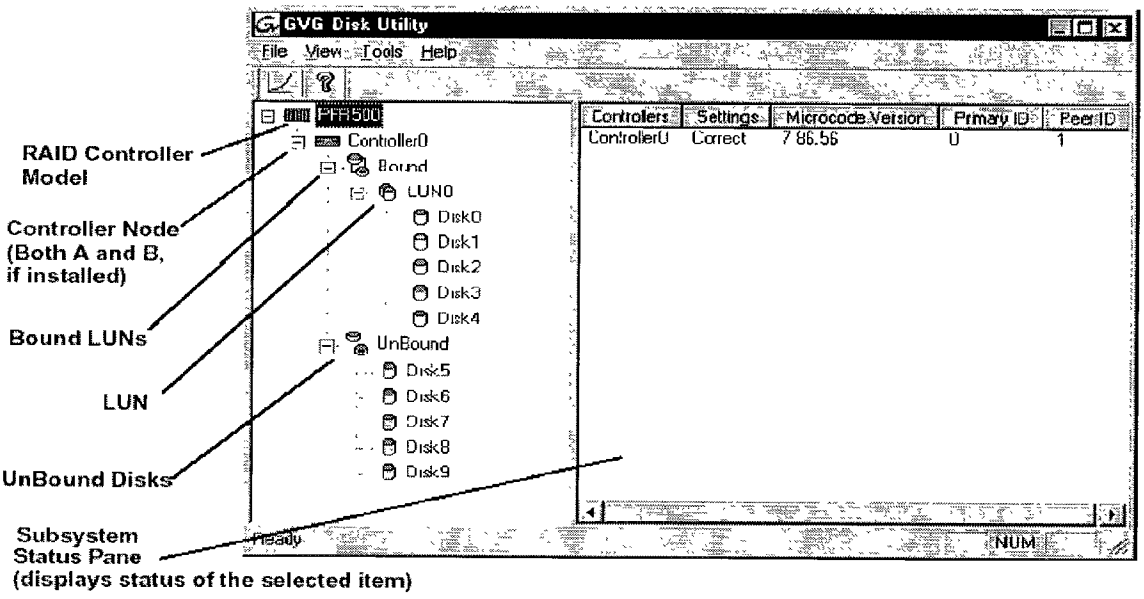
GVG Disk Utility
File View Tools Help
FFR500
Controller0
Bound
LUN0
Disk0
Disk1
Disk2
Disk3
Disk4
UnBound
Disk5
Disk6
Disk7
Disk8
Disk9
Controllers
Settings
Microcode Version
Primary ID
Peer ID
Controller0
Correct
7.86.56
0
1
Ready
NUM
RAID Controller Model
Controller Node (Both A and B, if installed)
Bound LUNs
LUN
UnBound Disks
Subsystem Status Pane (displays status of the selected item)

FIG. 1

Software Requirements Specification

For Disk Utility++ (NT based disk utility)

1 Introduction

1.1 Purpose

This document provides the software requirements specification for the NT based Disk Utility which will be common to both SANs, as well as stand-alone Profiles.

1.2 Scope

This document lays out the general requirements of the new disk utility. It is outside of the scope of this document to detail the vendor specific RAID configuration API. The vendor specific commands can be found in documentation provided to GVG from the RAID vendor.

1.3 Definitions, Acronyms, and Abbreviations

Profile	Profile XP or any future Profile
MFS	Media File System
CVFS	Centra Vision File System (SAN file system from ADIC)
SAN	Storage Area Network
fabric	The FC switch interconnects between the disks, servers, and clients
stand-alone	non-SAN
MAN	Media Area Network (ie: a SAN)
File System Server (FSS)	Dell server used as meta-data and data-base server for the SAN
RTS	Real Time System
HBA	Host Bus Adapter
FSS HBA	Fibre Channel HBA card that runs in the Dell File System Server
FCCard	Dual Fibre Channel Disk Card that runs vxWorks on Profile RTS
gvgscsi	NT-SCSI driver which interfaces NT apps with RTS-SCSI driver
RTS-SCSI driver	GVG RTS driver which runs on the FCCard
RAID	Redundant Array of Independent Disks
JBOD	Just a Bunch Of Disks (non-RAID)

RAID vendor	RAID vendor providing systems to GVG for use with the Profile
drive	For this document "drive" will refer to a single physical drive, rather than a LUN
LUN	Logical Unit Number (a group of physical drives represented as a single logical drive)
RAID3	RAID level which has a group of data drives, and a single dedicated parity drive
RAID1	mirrored pair of physical drives
Bind	To group together a set of physical drives into a LUN
Unbind	To destroy (or ungroup) the LUN
Controller Firmware	Microcode residing on the RAID controller
Drive firmware	Microcode residing on each individual drive
drive LEDs	Typically an activity LED and a status LED per drive
EMC/Clariion	Current RAID vendor
NEC	Future RAID vendor

2 General Description

2.1 Description

The current disk utility has a GUI which operates on NT that interfaces over KIPM to the RTS side. The actual configuration commands issued to the RAID controller are issued from the RTS-SCSI driver on the Profile. This prevents our current disk utility from being used on the Dell File System Server for the SAN. There are also other limitations with the current disk utility with regards to configuring varying LUN sizes and types and supporting multiple RAID vendors. While the current disk utility has served well its purpose, at this point it is reasonable to develop the new requirements into an NT based disk utility. The alternative is to develop the new requirements into the current disk utility, which will always be confined to run on a Profile.

The new disk utility will be an NT based application that can interface with any NT-SCSI driver (ie: gvgscsi when running on an XP, or a standard NT-SCSI driver when running on the Dell FSS).

2.2 General Constraints

The new disk utility should have no interaction with the RTS.

Currently for MFS, the file system configuration is combined with disk configuration. These will become separate entities, where the new disk utility does only that--disk configuration.

The File System configuration will be its own utility. One for the MFS, and one for the SAN (CVFS utilities). There will be no major changes to the current method of creating and naming a MFS File System. That portion of the GUI will simply become its own utility.

3 Functional Requirements

3.1 General Requirements

- Disk Utility must function on both a Profile or any Windows NT system that is attached to the RAID storage devices (either directly or over a FC fabric). At this time, there is not a requirement to support remote configuration.
- There must be a single disk utility for both SAN and non-SAN and all current and future Profiles.

3.2 SAN Specific Requirements

The ability to configure additional RAIDs on the SAN without interruption to RAIDs already in operation. This is not a requirement for stand-alone systems.

3.3 Stand-alone Specific Requirements

Disk Utility may not be launched if other Profile Applications are open.

3.3 Multiple RAID Vendor Requirements

- Must support any RAID vendor that is qualified to be used with a Profile (currently, that is only EMC/Clariion. Soon it will include NEC and should be part of the initial effort. Any subsequent RAID vendor can be phased in as needed).
- Since the specific RAID configuration API will vary from RAID vendor to vendor, each new RAID vendor should be represented as a class with its own methods that correspond to each required operation (such as bind a LUN--see below).

3.4 Disk Utility

- The ability to "discover" all devices including RAID controllers, drives and/or LUNs.
- A "Tree Control" existing of a list of all RAID controllers and their associated drives and/or LUNs. RAID controllers should list the vendor type and FC ID. Drives should list the drive # or ID, vendor type, model.
- The ability to select a group of drives, and bind them into either a RAID3 or RAID1 LUN. A LUN can consist of any number of drives (within vendor specific limits). There is no requirement to provide a "status" bar indicating rebuild progress. Per Bill S., there are issues preventing one from doing this from an NT application.
- Select a single LUN or group of LUNs, and unbind them.
- "unbind all LUNs"
- Select any number of drives and/or LUNs and blink the drive status LEDs. This allows one to physically look at and identify the drive or LUN. This feature may not be supported by certain RAID vendors.
- Select any single drive or group of drives and designate them as Hot Spare drives.
- Select any Hot Spare drive or group of Hot Spare drives and clear that designation.
- A button that allows the RAID controller clock to be reset to match the NT system clock (may be vendor specific).
- Select a single RAID controller, and load firmware. (There is no requirement to load firmware to *all* controllers at once--we don't want to give this capability since there may be mixed vendors). If LUNs are binding, fail this operation.
- Select a drive or group of drives and load drive firmware.
  If LUNs are binding, fail this operation.
- The ability to perform multiple of the above steps with only a single reboot at the end.

4 Other Considerations

Currently there exists code on both the NT and RTS that may be leveraged in this effort. Portions of the current disk utility GUI may be re-used. The major addition to the GUI would be the tree control with the associated logic to handle varying LUN sizes and RAID types. Much of the RTS code that issues the actual RAID configuration commands can be ported to NT. The work required to support multiple RAID vendors is considerable (ie: it won't be easy).

5 Version

Date	Revision	Description	By
April 11, 2001	V1.0	Initial Draft	Dave Crowther
May 2, 2001	V2.0	Added description for Hot Spare drives. Added "single reboot" requirement Clarified the requirement for running on any NT system.	D.C.

Chapter 3

Working with Profile XP storage systems

This chapter introduces the Fibre Channel RAID storage system and describes how to connect and configure storage on the Profile XP system. Topics included are:

PFC500 product description

The PFC500 is a 3.5 RU 10 drive Fibre Channel disk array with RAID drives, hot spare drives, dual power supplies (optional), redundant controllers (optional) and spare fans – all hot swappable. The PFC500 comes with either 36 GB (PFC536) or 73 GB drives (PFC573).

PFC500 expansion is available using the PFC500E Expansion Chassis which houses 5 or 10 drives, but does not need a RAID controller. The PFC500's RAID controller will control up to 2 expansion frames (30 drives total). All PFC500 chassis must contain the same size disk modules to make full use of the storage capacity.

> ***NOTE: Grass Valley Group does not support mixing disk drives of differing capacities in any RAID chassis connected to a Profile XP media platform. All disk drives in any RAID chassis connected to a Profile XP media platform must be of the same capacity. For example, if a PFC500 and a PFC500E are connected to a Profile XP system, all the disk drives in both the PFC500 and the PFC500E must be of the same capacity.***

The PFC500 offers optional redundant RAID controllers for automatic failover. Failover can occur when the primary RAID Controller, Fibre Channel Port of the Profile XP, or Fibre Channel cable fails. When failover occurs, the system continues to run with minimal interruption.

For additional information on the PFC500 Fibre Channel RAID chassis and the PFC500E Expansion chassis, refer to the *PFC500 Instruction Manual.*

PFR500 product description

The PFR 500 contains five or ten half-height 3.5" Fibre Channel Arbitrated Loop (FC-AL) disk drives. The chassis also supports one or two hardware RAID Controllers in one 3U high rack-mountable chassis. The PFR 500 currently uses 73GB or 180GB drive capacities. With ten drives, one chassis holds up to 730GB or 1.8TB depending on disk drive option.

The PFR 500E RAID Expansion Chassis provides additional storage capacity. It is an identical chassis with two Loop Bypass Board installed. Up to nine PFR500E RAID Expansion Chassis can be connected to a single PFR500 comprising a single disk-array storage system with a total of 100 drives and 7.3TB or 18TB of storage depending on disk drive option. The built-in chassis daisy-chaining capabilities provide for cost effective storage expansion as requirements grow.

The PFR 500 utilizes dual FC-AL technology, allowing two loop configurations within a single chassis. Port-Bypass Circuits have been added to maintain loop integrity during failures without user intervention. Each loop and associated Port Bypass Circuits along with all other active components are on redundant, separate hot swappable modules. This improves serviceability and increases fault tolerance by eliminating any single point of failure. With two RAID Controllers, the two loops within a single standard chassis are configured as a single loop with a backup loop in standby mode.

PFR500 host interface supports both Arbitrated Loop and fabric switching industry standard Fibre Channel technology. The physical layer is determined by the use of GBICs.

NOTE: Grass Valley Group does not support mixing disk drives of differing capacities in any RAID chassis connected to a Profile XP media platform. All disk drives in any RAID chassis connected to a Profile XP media platform must be of the same capacity. For example, if a PFR500 and PFR500E are connected to a Profile XP system, all the disk drives in both the PFR500 and the PFR500E must be of the same capacity.

Storage capacity estimates

The storage capacity is determined by the number and size of disk drives in system, and the video data rate selected to record video. The Fibre Channel RAID storage system can be configured with five to thirty disk drives, including the Expansion chassis.

You can refer the following tables for storage capacity estimates based on the number of disk drives installed and the video data rate. The estimates in these tables are based on use of four audio channels and one time code channel per video channel using the 525 video format.

Storage capacity estimates in hours using 36 GB drives:

Video Data Rate	Number of Drives					
	5	10	15	20	25	30
4 Mbs	37.8	75.6	113.4	151.2	189	226.8
8 Mbs	25.4	50.8	76.2	101.6	127	152.4
12 Mbs	19 1	38.2	57.3	76.4	95.5	114.6
15Mbs	16 1	32 2	48 3	64.4	80.5	96.6
18 Mbs	13 8	27 6	41 4	55.2	69	82.8
24 Mbs	10.8	21.6	32 4	43.2	54	64 8
32 Mbs	8.4	16.8	25.2	33.6	42	50.4
40 Mbs	6 8	13 6	20 4	27 2	34	40.8
50 Mbs	5.5	11	16 5	22	27 5	33

Storage capacity estimates in hours using 73 GB drives

Video Data Rate	Number of Drives					
	5	10	15	20	25	30
4 Mb/s	76 6	153 2	229.9	306 4	383.1	459 6
8 Mb/s	51 4	102 8	154.4	205.6	257.3	308 4
12 Mb/s	38.7	77 4	116.1	154 8	193.5	232.2
15 Mb/s	32.6	65.3	97.9	130.5	163.1	195.7
24 Mb/s	22.2	44 3	66.4	88.6	110.7	132.9
30 Mb/s	18.2	36.5	54.7	72.9	91.1	109.4
40 Mb/s	14.1	28.1	42.2	56.2	70 3	84.3
50 Mb/s	11 4	22 9	34 3	45 7	57.1	68 5
65 Mb/s	8 8	17.6	26 3	35 1	43.9	52.6

Maximum video data rate per channel for PFC500

The maximum video data rate per channel depends on the Profile XP model, the number of PFC500 chassis connected, and the number of disk drives in the system.

The Profile XP system accesses the disk array through the PFC500' RAID Controller board. As video data rates and number of channels in operation increase, the bandwidth of a single RAID controller board may be exceeded. Some Profile XP models require at least two or three PFC500 to satisfy the bandwidth required.

The following table shows the minimum storage configuration required for your Profile XP model for different video data rates using the PFC500 Fibre Channel RAID Chassis. Please consult your Grass Valley Group representative for minimum storage configuration information using the PFR500.

Profile XP Model	Number of Channels	Video Data Rate Per Channel[a]					
		8Mbs	12Mbs	15Mbs	24Mbs	30Mbs	50Mbs
PVS1022	4	5 drives 1 PFC500	5 drives 1 PFC500	5 drives 1 PFC500	10 drives 1 PFC500	15 drives 1 PFC500	20 drives 2 PFC500
PVS1024	6	5 drives 1 PFC500	10 drives 1 PFC500	10 drives 1 PFC500	15 drives 1 PFC500	20 drives 1 PFC500	30 drives 3 PFC500
PVS1044 PVS1026 PVS1062 PVS1008	8	10 drives 1 PFC500	10 drives 1 PFC500	10 drives 1 PFC500	20 drives 1 PFC500	20 drives 2 PFC500	30 drives 3 PFC500

[a]2 channels audio and 1 timecode channel. Does not take into account video network transfers. PFC500 refers to PFC518 or PFC536.

You can select the video data rate the Profile XP system uses to record the video signal. When selecting video data rates, be sure not to exceed the total maximum system bandwidth. Fibre channel transfers, archive operations, and high bandwidth recording or playout may require that you reduce the data rate on other channels or operate fewer channels at one time. Although playout and recording take priority over other operations, exceeding your Profile XP system bandwidth may result in record or playback problems. For information on setting the video data rate, refer to "Selecting video quality for a channel" on page 126.

Connecting the RAID storage chassis

The *Installation Guide* for your PVS Series model provides complete information on connecting your RAID storage chassis to the Profile XP system using both single and redundant RAID controllers. It also describes how to set Fibre Channel Loop address and Chassis ID for the RAID storage and the Expansion Chassis.

Configuring storage using GVG Disk Utility

After connecting the Fibre Channel RAID storage system to the Profile XP Media Platform, you must use the GVG Disk Utility to bind the storage system's disk modules into logical units (LUNs) and create a video file system that is recognizable by the Profile XP Media Platform.

This section includes the following:

About configuring storage

This section describes terms and concepts related to configuring your the Fibre Channel RAID storage system. Topics include:

- Binding disk modules
- Unbinding disk modules
- Creating a video file system
- Expanding Storage
- File system volume name
- Using Hot Spare disks

Binding disk modules

Binding the disk modules hardware formats them into a logical units called LUNs. The disks that make up a LUN are accessed as a contiguous disk space. Disk modules must be bound before they can be part of the video storage file system.

You can use the GVG Disk Utility to bind disks modules as Data LUNs (RAID-3 containing 4 data disks and 1 parity disk) or as hot spare LUNs, depending on the storage needs of your system. Depending on the model, a RAID chassis can contain up to 10 or 15 drives.

For simplicity, the disk utility only allows binding contiguous disk modules into LUNs, that is, in a ten disk chassis, disks 0-4 and 5-9 must be bound as two different LUNs. To help in finding the physical location of disks in large systems, the disk utility includes an **Identify Disk** button which allows the selected disk drive's disk access light to flash.

After binding, disk modules become slot specific and cannot be moved to other disk module slots.

Unbinding disk modules

Unbinding reverses the bind process. Unbinding might be needed when reconfiguring a storage system, for example, when reconfiguring hot spare disks into a data LUN. PFC500 storage systems allow unbinding LUNs independently. PFR500 storage systems inforce unbinding all LUNs simultaneously (i.e. clearing the configuration).

>>> ***CAUTION: Unbinding always destroys all user media files stored on the disk modules.***

Creating a video file system

The Fibre Channel RAID storage system must contain a video file system recognizable by the Profile XP Media Platform system. When you choose to make a new video file system, the Disk Utility creates one video file system using all available data LUNs. All hot spare LUNs and Unbound LUNs are ignored.

>>> ***CAUTION: Creating a video file system destroys the existing file system and all user media files stored on the disk modules.***

Expanding Storage

You can expand storage by installing additional disk modules in empty disk slots or by connecting additional Fibre Channel RAID Chassis or Expansion chassis. To use the additional storage, you must bind the new disk modules and create a new video file system.

The disk modules you use in a LUN must all have the same capacity to fully use the modules' disk space. The same applies to LUNs. Make sure all LUNs have the same capacity to fully use the LUN's disk space.

File system volume name

The video file system volume name is the logical name used when accessing the storage system. All local and remote Profile XP applications use the volume name when referencing the video file system.

You can change the video file system volume name at any time without affecting the media files stored on the system. Changing the volume name may be required with some third party Profile XP control software which requires a specific volume name.

Using Hot Spare disks

You can bind disks as hot spare drives. Hot spare drive are available for automatic failover in the event a disk module fails in a data LUN. When the failed disk module is replaced, the RAID controller copies the hot spare drive data to the replacement drive, then returns the hot spare drive back to hot spare status.

If the drives you want to designate as hot spares are bound as part of a data LUN, you must unbind the drives first, then bind them as hot spares. Hot spare drives are available only to the RAID Controller in the Fibre Channel RAID Chassis controlling them.

Introducing the GVG Disk Utility

This section includes the following:

- GVG Disk Utility user interface overview
- Checking storage subsystem status

GVG Disk Utility user interface overview

The GVG Disk Utility user interface includes a tree view in the left-hand pane, and a status information area displayed in the right-hand pane. The tree view displays the hardware that make up the storage system connected. The context menus in the tree view are used to configure storage. The right-hand status pane displays information about the item selected in the tree view.

Hierarchy of the tree view:

RAID Controller Model - Provides a logical grouping of RAID chassis models, i.e. all PFC500 RAID Controller Chassis connected are listed under the PFC500 icon.

Controller - Represents the RAID Controllers found. These are numbered in the order discovered. The controller icon represents both RAID Controller A and, if installed, RAID Controller B. To determine if an optional RAID Controller B is installed, select the Controller icon in the tree view, then examine the status pane.

Bound LUNs - Expanding the Bound node displays all bound LUNs.

LUN - Represents a bound LUN. Expanding the LUN node displays the disk modules that make up the LUN.

UnBound LUNs - Expanding the UnBound node, displays all unbound disk modules.

Disks - Represents the disk modules.

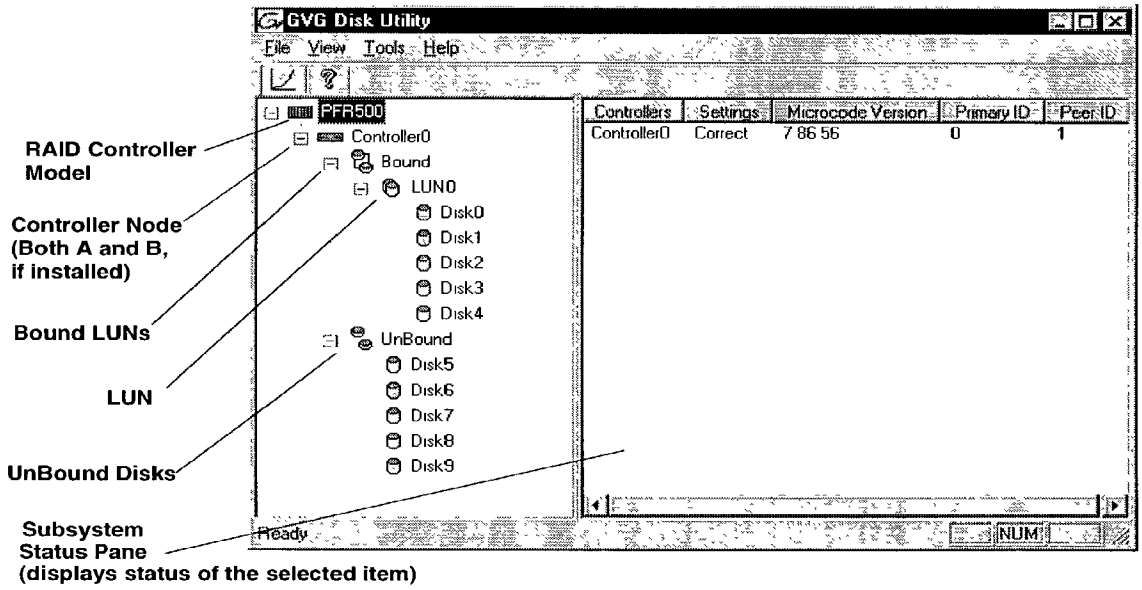

Checking storage subsystem status

NetCentral monitoring software is available to monitor your Profile XP and the attached RAID storage system, however, some limited status information for storage subsystems is displayed in the GVG Disk Utility. This can be helpful when configuring storage.

You can view status information by selecting these items in the tree view.

- Storage Model
- Controller
- Bound
- LUN
- Disk

Binding disks and creating a video file system

Before you can use your PVS 1000 series system, you must create a video file system on the Fibre Channel RAID Chassis using the GVG Disk Utility.

Creating a video file system involves:

- Setting the RAID Controller Fibre Channel ID (PFR 500s only)
- Binding any unbound disk modules
- Creating and naming the file system

NOTE: If you want to use hot spare drives in your system, refer to the procedures in "Configuring hot spare drives" on page 74.

To create a video file system:

1. Start Configuration Manager using the desktop shortcut or by selecting **Start | Programs | Profile Applications | Configuration Manager**. The Configuration Manager dialog box appears.

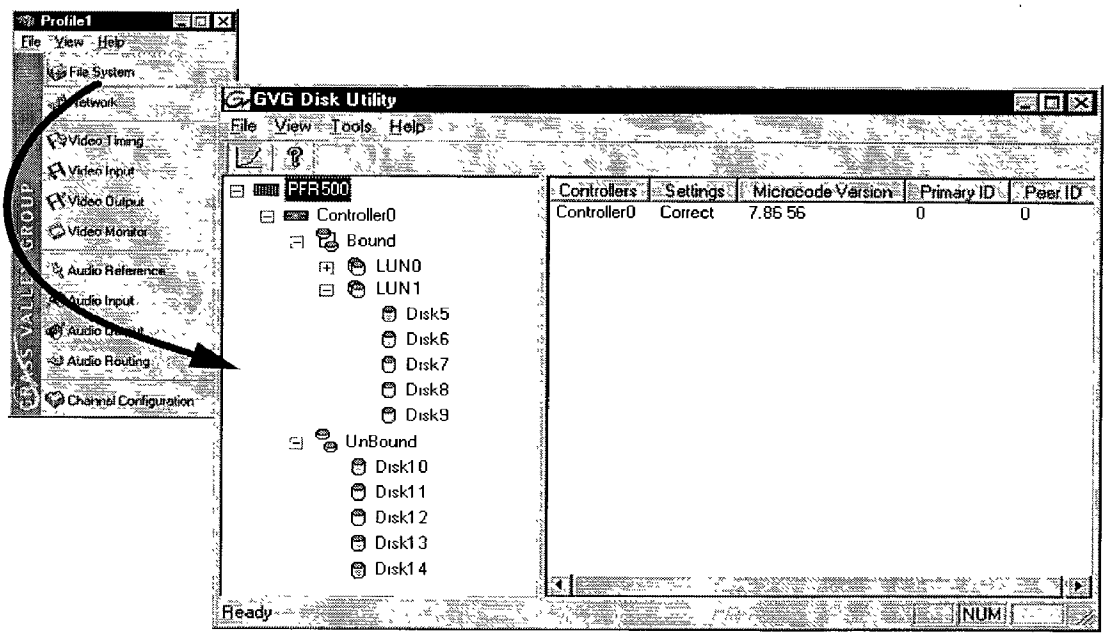

NOTE: For simplicity, this example shows only a 15 disk system.

2. Click **File System** to start the GVG Disk Utility.
3. If you are configuring a PFR 500 storage system, perform the following steps to set the Fibre Channel ID, otherwise, proceed to step 4.
   a. Expand the tree view and right-click the **Controller0** icon under PFR 500, then

select **Set Fibre Channel ID** in the context menu.

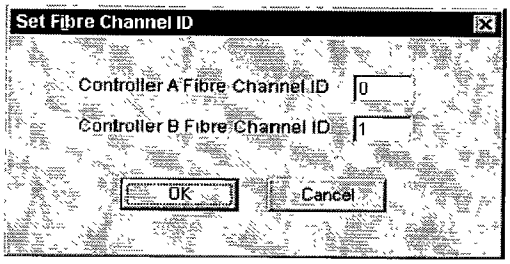

b. In the Set Fibre Channel ID dialog box, enter IDs as shown, then click **OK**.

c. Click **OK** in the Operation Successful dialog box.

d. Cycle the power on the PFR500 chassis containing the RAID controllers.

e. After the PFR500 has finished initializing, reboot the Profile XP.

f. Start Configuration Manager using the desktop shortcut or by selecting **Start | Programs | Profile Applications | Configuration Manager**. The Configuration Manager dialog box appears.

g. Click **File System** to start the GVG Disk Utility.

4. Expand the tree view and identify bound LUNs and unbound disks by their placement in the hierarchy of the tree view.

5. If there are unbound LUNs displayed in the tree view, you must perform the following steps, otherwise, proceed to step 6.

   a. To bind unbound disks, right-click the **Unbound** node for a controller, then select **Bind LUNs** in the context menu. (Peer controllers that share the same set of disks are automatically selected as a pair.)

   The Binding LUN... dialog box opens showing all unbound disks for the controller listed in the Available Disk list.

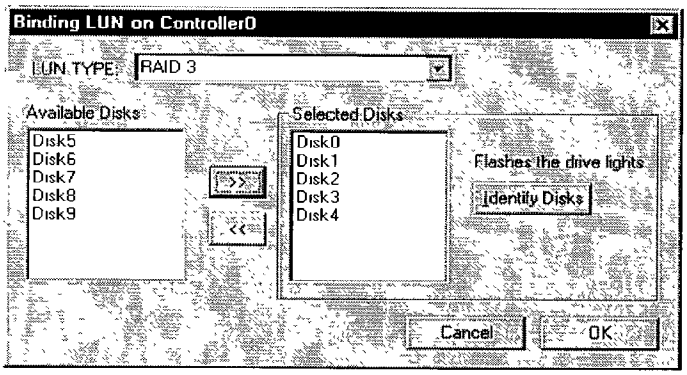

b. Select RAID 3 using the LUN TYPE drop-down box.

c. Select five contiguous disks, then click the add button [>>] to add them to the Selected Disks list. (TIP: Use 'shift-click' or 'control-click' to select disks.)

NOTE: As an aid in identifying a disk module's physical location, select it in the Selected Disks list, then click* Identify Disks. *This causes the disk drive LED to flash.

d. Click **OK** to close the Binding LUN... dialog box and begin the binding process. The Progress Report dialog box opens, showing the status of the binding process.

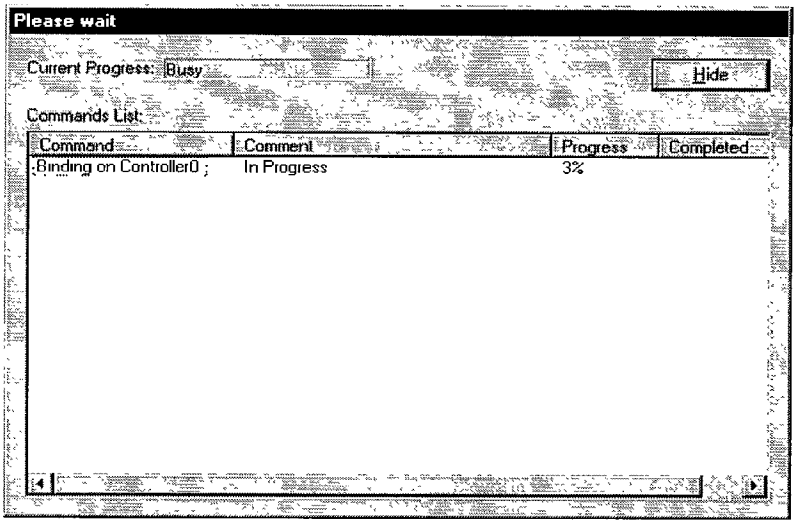

e. Repeat step a through step d for all unbound disks on all remaining controllers displayed in the tree view.

NOTE: PFR500 controllers take less than a minute to bind a LUN, but can bind only one LUN at a time. You must wait until binding is completed before binding the next LUN. PFC500 controllers take more than an hour to bind a LUN, but allow multiple binding processes at the same time.

f. Upon 100% completion, click **Close** in Progress Report window.

NOTE: Do not proceed until all LUNs are finished binding.

g. Reboot the Profile XP System.

h. Start Configuration Manager using the desktop shortcut or by selecting **Start | Programs | Profile Applications | Configuration Manager.** The Configuration Manager dialog box appears.

i. Click **File System** to start the GVG Disk Utility.

6. In the GVG Disk Utility main menu, click **Tools | Make File System**. The File System wizard appears.

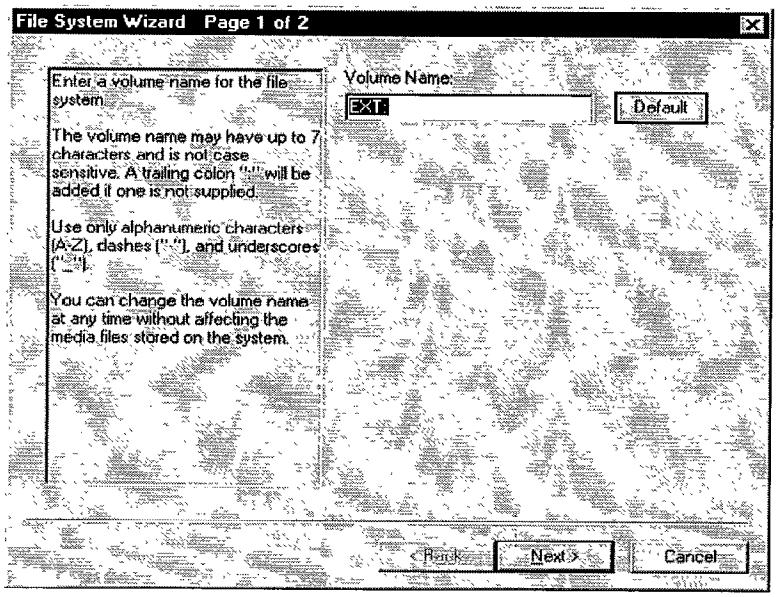

7. Enter a volume name or accept the default name, then click **Next**.

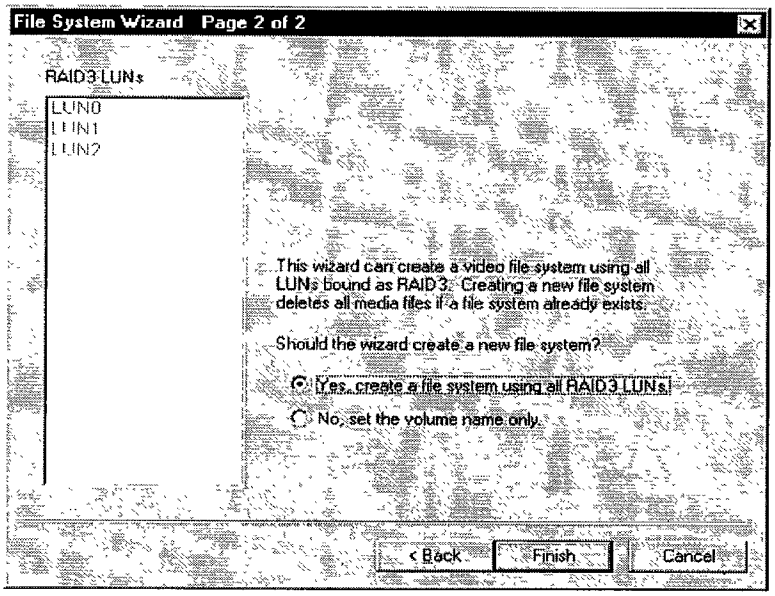

8. Select **Yes, create a file system using all RAID3 LUNs**, then click **Finish**.
9. Click **OK**, in the warning message box. The Progress Report window appears showing the create file system task status.

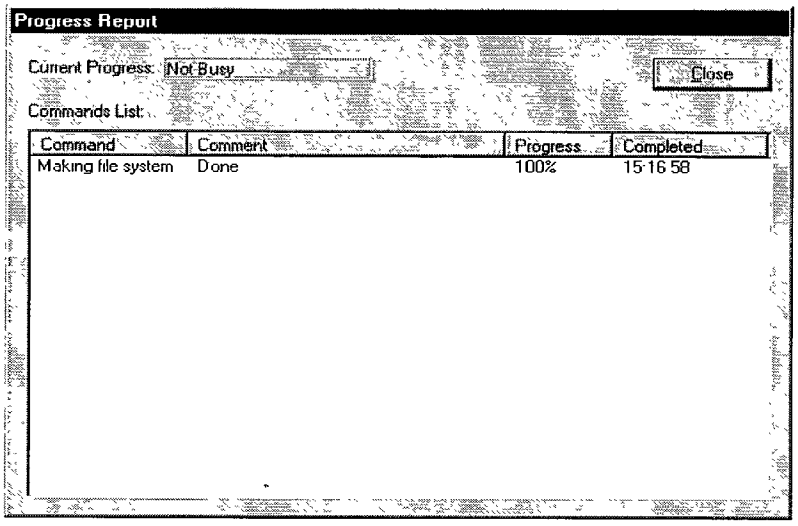

10. Upon 100% completion, click **Close** in Progress Report window.
11. When the reboot message appears, click **OK** to reboot the Profile XP system. Afterward, the Fibre Channel RAID Storage System is ready for operation.

Configuring hot spare drives

When you configure storage, you can designate drives as hot spare drives. Hot spare drives are used in the event of a drive failure in a data LUN. If a drive failure occurs, the RAID Controller automatically selects a hot spare drive to use in place of the failed drive. When the failed disk module is replaced, the RAID controller copies the hot spare drive data to the replacement drive, then returns the hot spare drive back to hot spare status.

If the drives you want to designate as hot spares are bound as part of a RAID-3 LUN, you must unbind the drives first, then bind them as hot spares. (Refer to "Unbinding LUNs for reconfiguration" on page 83.) You can configure any disk in the array as a hot spare drive, however, hot spare drives are available only to RAID Controller to which they are connected.

Perform the following procedure if you are configuring Fibre Channel RAID storage systems to include hot spare drives.

To create hot spare drives:

1. Start Configuration Manager using the desktop shortcut or by selecting **Start | Programs | Profile Applications | Configuration Manager**. The Configuration Manager dialog box appears.

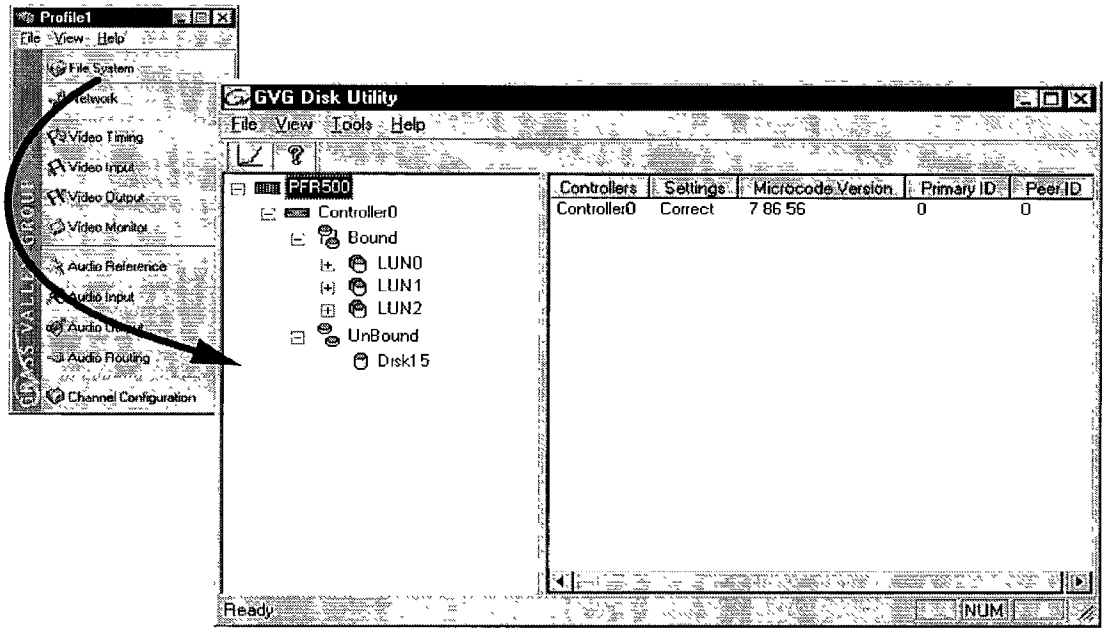

2. Click **File System** in the Configuration Manager window. The GVG Disk Utility appears.

NOTE: For simplicity, this example shows only a 16 disk system.

3. To bind unbound disks as hot spares, right-click the **Unbound** node for a controller,

then select **Bind LUNs** in the context menu. (Peer controllers that share the same set of disks are automatically selected as a pair.)

The Binding LUN... dialog box opens showing all unbound disks for the controller listed in the Available Disk list.

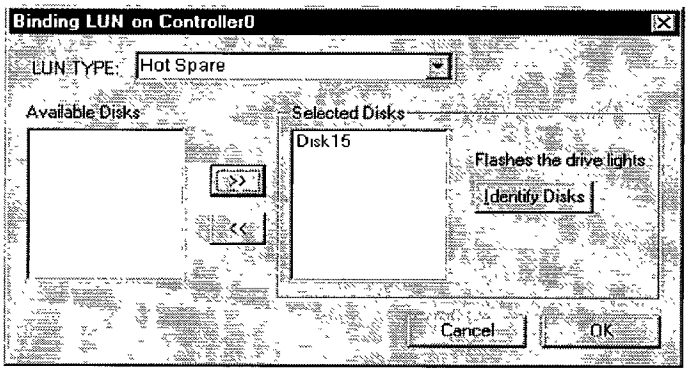

4. Select Hot Spare using the LUN TYPE drop-down box.
5. Select the disk(s) to be used as hot spares, then click the add button >> to add them to the Selected Disks list.

   NOTE: As an aid in identifying a disk module's physical location, select it in the Selected Disks list, then click* Identify Disks. *This causes the disk drive LED to flash.

6. Click **OK** to close the Binding LUN... dialog box and begin the binding process. The Progress Report dialog box opens, showing the status of the binding process.

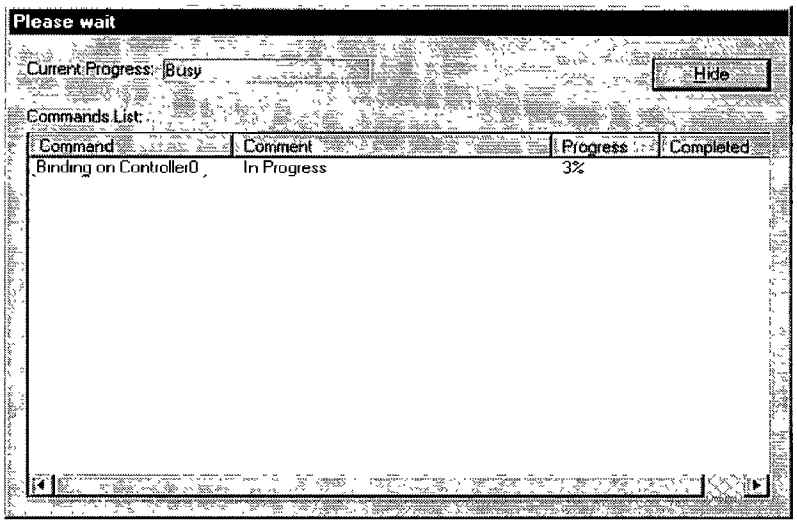

7. Repeat step 3 through step 6 to bind additional hot spare disks modules on other controllers as needed.
8. Upon 100% completion, click **Close** in Progress Report window.
9. Reboot the Profile XP System.
10. If you have not created a video file system on the remaining disks, refer to "Binding disks and creating a video file system" on page 69.

Changing the video file system volume name

You can change the video file system volume name at any time without affecting the media files stored on the system. Changing the volume name may be required with some third party Profile XP control software which requires a specific volume name.

To change the file system volume name:

1. Start Configuration Manager using the desktop shortcut or by selecting **Start | Programs | Profile Applications | Configuration Manager**. The Configuration Manager dialog box appears.

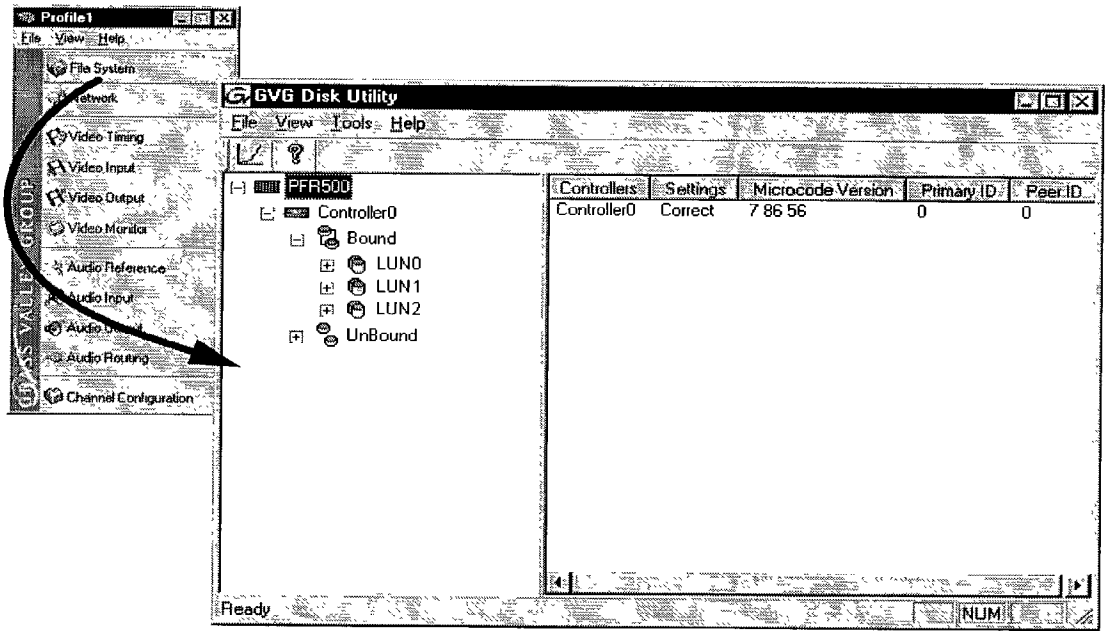

2. In the GVG Disk Utility main menu, click **Tools | Make File System**. The File System wizard appears.

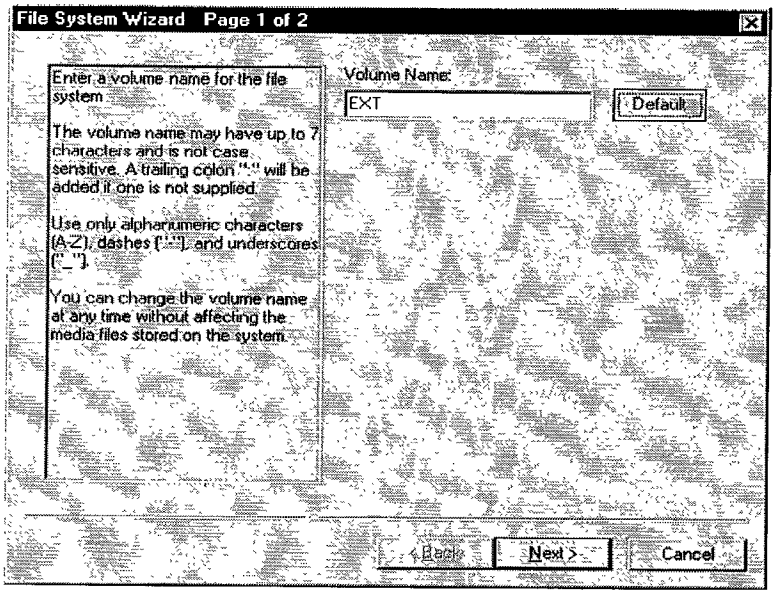

3. Enter a new volume name, then click **Next**.

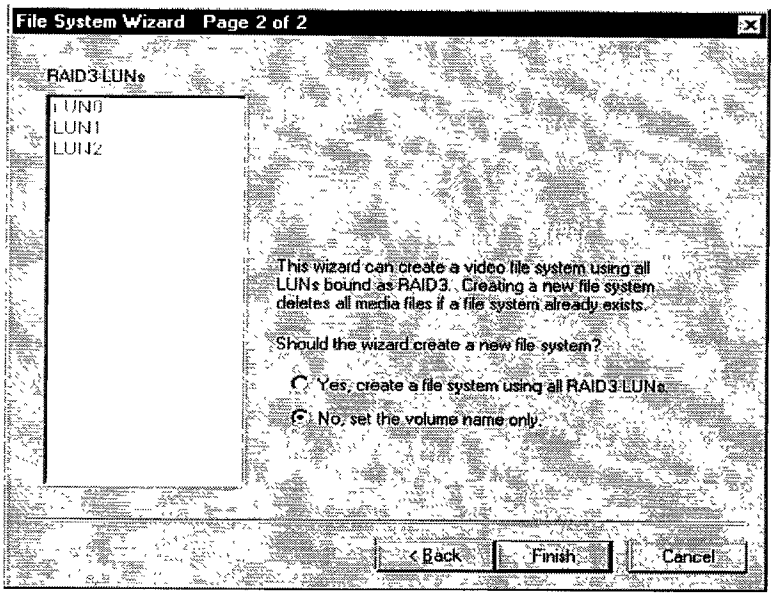

4. Select **No, set the volume name only**, then click **Finish**. The Progress Report window appears showing the change volume name task status.
5. Upon 100% completion, click **Close** in Progress Report window. Be sure to restart applications that access disk storage on the Profile XP system.

Expanding storage

You can expand storage by installing additional disk modules in empty disk slots or by connecting additional storage chassis. To use the additional storage, you must bind unbound LUNs and create a new video file system which includes the new disk modules.

>>> ***CAUTION: Back up all media on existing drives before creating a new video file system. Creating a new video file system destroys all media files and other data on the disk modules.***

If you are adding additional storage chassis, refer to the *Installation Guide* for your Profile XP Media Platform to make hardware connections, then return to this procedure to configure the storage system.

To create a video file system on expanded storage:

1. Start Configuration Manager using the desktop shortcut or by selecting **Start | Programs | Profile Applications | Configuration Manager**. The Configuration Manager dialog box appears.

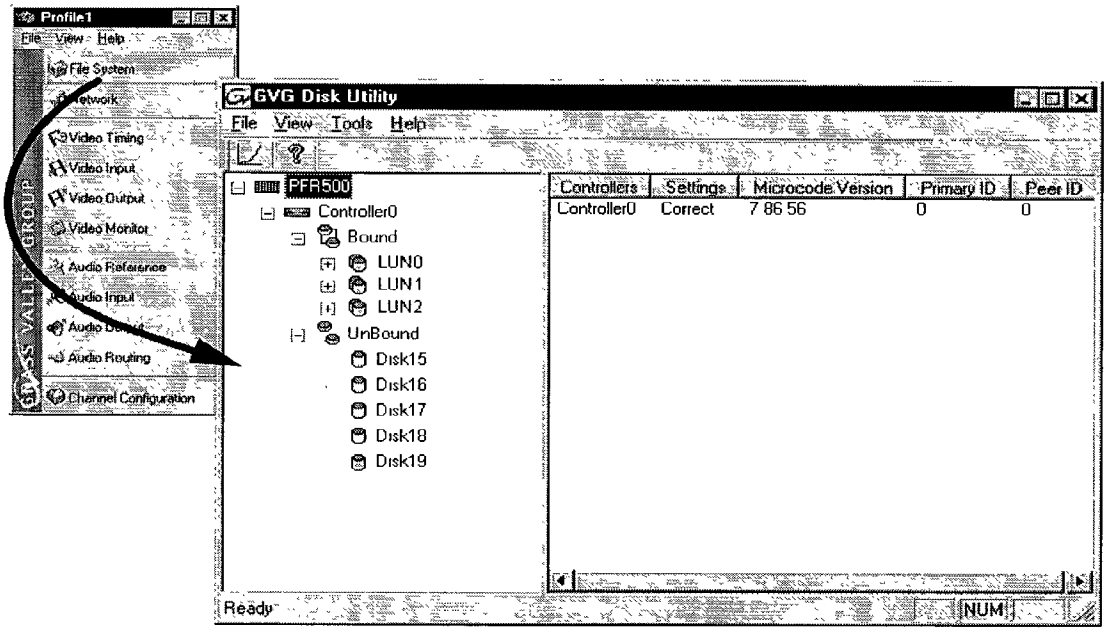

2. Click **File System** in the Configuration Manager window. The GVG Disk Utility appears.

   NOTE: For simplicity, this example shows only a 20 disk system in the Disk Utility.

3. Expand the tree view and identify bound LUNs and unbound disks by their placement in the hierarchy of the tree view.

4. If there are unbound LUNs displayed in the tree view, you must perform the following steps, otherwise, proceed to step 6.

   a. To bind unbound disks, right-click the **Unbound** node for a controller, then select **Bind LUNs** in the context menu. (Peer controllers that share the same set of disks are automatically selected as a pair.)

      The Binding LUN... dialog box opens showing all unbound disks for the controller listed in the Available Disk list.

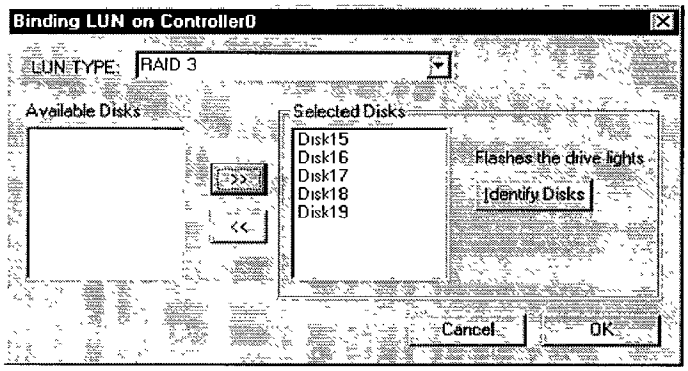

   b. Select RAID 3 using the LUN TYPE drop-down box.

   c. Select five contiguous disks, then click the add button [>>] to add them to the Selected Disks list. (TIP: Use 'shift-click' or 'control-click' to select disks.)

      NOTE: As an aid in identifying a disk module's physical location, select it in the Selected Disks list, then click* Identify Disks. *This causes the disk drive LED to flash.

   d. Click **OK** to close the Binding LUN... dialog box and begin the binding process. The Progress Report dialog box opens, showing the status of the binding process.

   e. Repeat step a through step d for all unbound disks on all remaining controllers displayed in the tree view.

      NOTE: PFR500 controllers take less than a minute to bind a LUN, but can bind only one LUN at a time. You must wait until binding is completed before binding the next LUN. PFC500 controllers take more than an hour to bind a LUN, but allow multiple binding processes at the same time.

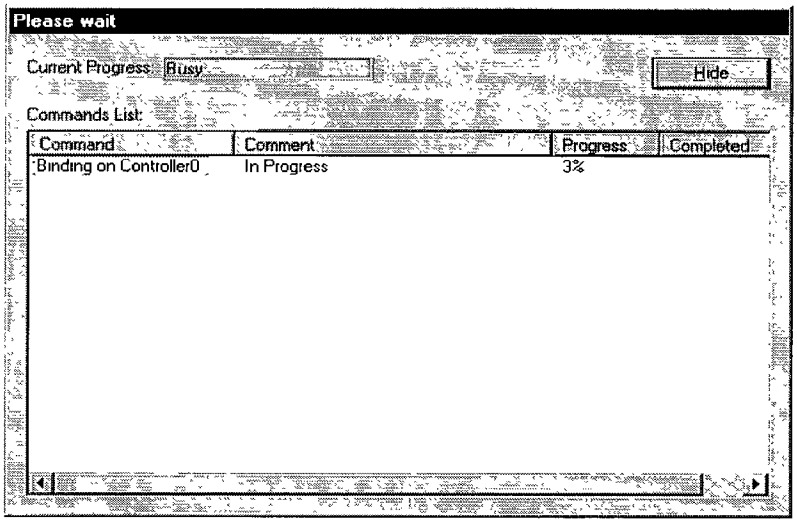

f. Upon 100% completion, click **Close** in Progress Report window.

NOTE: Do not proceed until all LUNs are finished binding.

g. Reboot the Profile XP System.

h. Start Configuration Manager using the desktop shortcut or by selecting **Start | Programs | Profile Applications | Configuration Manager**. The Configuration Manager dialog box appears.

i. Click **File System** to start the GVG Disk Utility.

5. In the GVG Disk Utility main menu, click **Tools | Make File System**. The File System wizard appears.

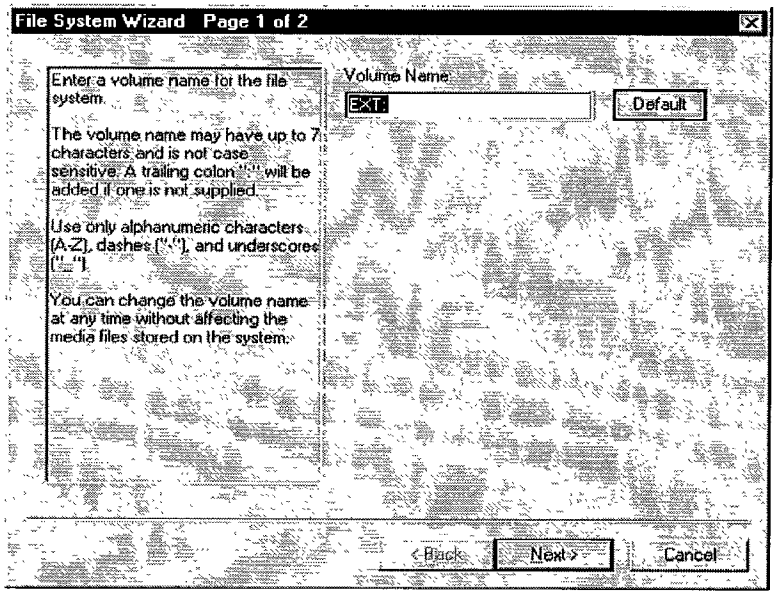

6. Enter a volume name or accept the default name, then click **Next**.

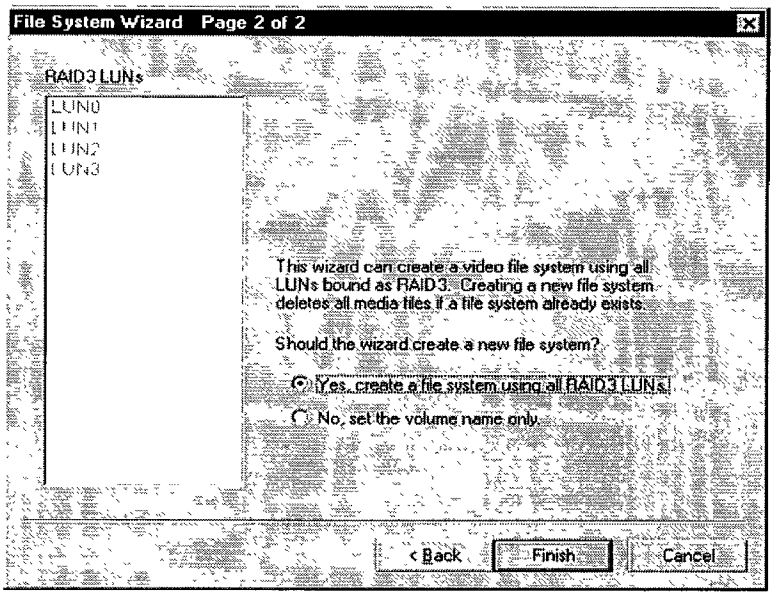

7. Select **Yes, create a file system using all RAID3 LUNs**, then click **Finish**.
8. Click **OK**, in the warning message box. The Progress Report window appears showing the create file system task status.

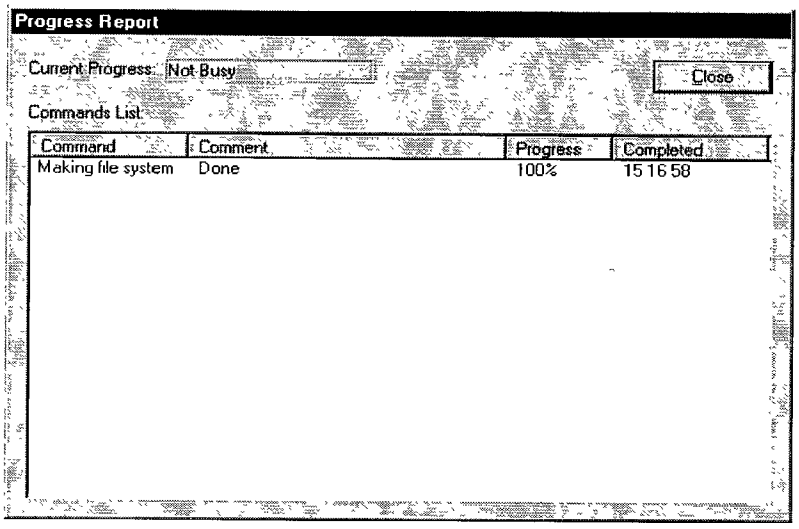

9. Upon 100% completion, click **Close** in Progress Report window.
10. When the reboot message appears, click **OK** to reboot the Profile XP system. Afterward, the Fibre Channel RAID Storage System is ready for operation.

Unbinding LUNs for reconfiguration

Unbinding may be necessary to reconfigure the disk modules in the RAID storage system. You must unbind a LUN before you can bind the disk modules in a different configuration.

>>> ***CAUTION: Unbinding destroys all user media files stored on the disk modules.***

This section includes the following procedures:

- Unbinding LUNs: PFC500 Storage Systems
- Unbinding LUNs: PFR500 Systems

Unbinding LUNs: PFC500 Storage Systems

If you need to unbind LUNs, perform the following:

1. Start Configuration Manager using the desktop shortcut or by selecting **Start | Programs | Profile Applications | Configuration Manager**. The Configuration Manager dialog box appears.
2. Click **File System** to start the GVG Disk Utility.
3. Expand the tree view and identify bound LUNs.
4. Right-click a LUN you want to unbind, then select **Unbind LUNs** in the context menu. The Unbinding LUN... dialog box opens.
5. On receiving the warning message, click **OK**.

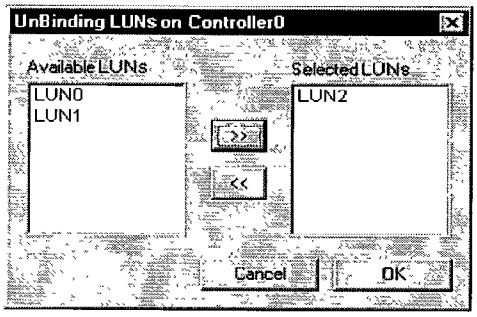

6. Move LUNs to the Selected LUNs box using the arrow buttons. (HINT: To select all LUNs, use **Shift + click** or **Shift + click**, then click the >> arrow button.).
7. Click **OK** to close the Unbinding LUN... dialog box and begin the unbinding process. The Progress Report dialog box opens, displaying a unbind status.
8. Repeat step 4 through step 7 to unbind LUNs on additional controllers, if needed.
9. Upon 100% completion, click **Close** in the Progress Report dialog box.

   This completes the unbind procedure.

Unbinding LUNs: PFR500 Systems

PFR500 RAID controllers do not allow unbinding individual LUNs. You must unbind all LUNs to reconfigure your storage system.

>>> ***CAUTION: In this software release, do not use the Unbind LUN entry in any context menu. You must use Unbind All LUNs in the Controller context menu.***

If you need to unbind LUNs, perform the following:

1. Start Configuration Manager using the desktop shortcut or by selecting **Start | Programs | Profile Applications | Configuration Manager**. The Configuration Manager dialog box appears.
2. Click **File System** to start the GVG Disk Utility.
3. Expand the tree view to identify bound LUNs.
4. Right-click the Controller icon, then select **Unbind All LUNs** in the context menu. The unbind all LUN operation begins.
5. On receiving the warning message, click **OK**.
6. On receiving the "Operation Successful" message, click **OK**.

   This completes the unbind procedure. You can now bind disk modules in a new configuration as required.

Identifying disk modules

The Disk Utility **Identify** feature allows you to the flash disk access LEDs so that you can physically locate a specific disk module. You can use identify features to physically locate a LUN, or individual disk.

To identify the disks in a LUN:

1. Right-click the LUN icon in the tree view, then select **Identify Disks** in the context menu.
2. Perform one of the following steps:
   - For PFR500 systems, click **OK** to identify disk modules in the selected LUN. Disk modules remain identified for 30 seconds. You can locate the disk modules now.

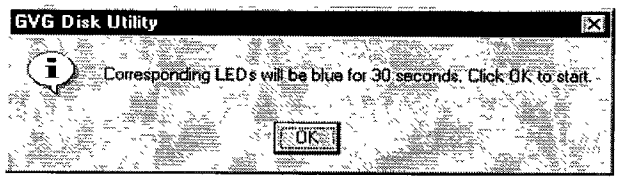

   - For PFC500 systems, locate the disk modules now by locating the flashing yellow LEDs. Click **OK** to extinguish the flashing LED.

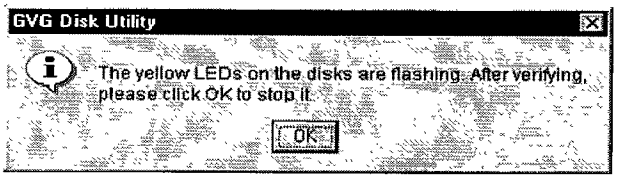

To identify a disk module:

1. Right-click the disk icon in the tree view, then select **Identify Disk** in the context menu.
2. Perform one of the following steps:
   - For PFR500 systems, click **OK** to identify disk modules in the selected LUN. Disk modules remain identified for 30 seconds. You can locate the disk modules now.

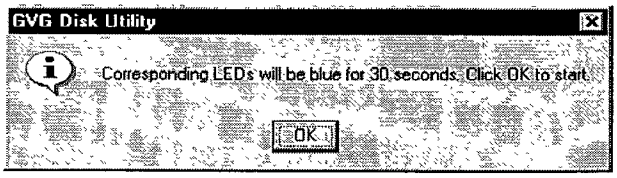

- For PFC500 systems, locate the disk modules now by locating the flashing yellow LEDs. Click **OK** to extinguish the flashing LED.

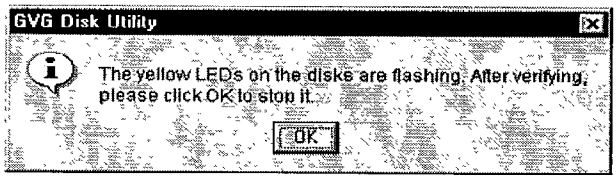

This concludes the disk identification procedure.

Storage maintenance tasks using GVG Disk Utility

This section includes maintenance procedures for your storage system. These procedures use the GVG Disk Utility to perform the following tasks:

Loading disk drive firmware

If you upgrade your Profile system software, you may be instructed in the software release notes to update the disk drive firmware on all the RAID chassis. This allows you to take advantage of the disk drive enhancements and benefit from improved reliability.

> ***NOTE: The disk drives on each controller are updated one at a time which can take as long as 2 minutes per drive. Disk Utility does allow multiple controllers to update simultaneously. Take this into consideration when scheduling the upgrade.***

If you need to update disk drive firmware, do the following:

1. In the GVG Disk Utility, right-click a controller in the tree view, then select **Advanced | Download Disk Firmware** in the context menu.
2. In the Open dialog box, locate and select the file containing the firmware version number indicated by the release notes. (Default microcode location on Profile XP systems is *c:\profile\microcode* directory)
3. Click **Open.** The Progress Report window appears showing the disk firmware update task and the percentage completion.
4. Perform step 1 through step 3 to update disk drives for remaining controllers.
5. Reboot all RAID storage chassis and expansion chassis.
6. Reboot the Profile XP system.

   This concludes the Loading Disk Firmware procedure.

Loading RAID controller firmware

If you upgrade your Profile system software, you may be instructed in the software release notes to update the RAID Controller microcode on all the RAID chassis. This allows you to take advantage of the RAID enhancements and benefit from improved reliability.

If you need to update your microcode, do the following:

1. In the GVG Disk Utility, do one of the following
   - To update a single controller, select the controller in the tree view. Redundant controllers that share the same set of disks are automatically selected and upgraded as a pair.
   - To update all controllers, select the RAID model node in the tree view, such as the PFC500 or PFR500 icons, to select all the controllers under the node.
2. Click **File | Load microcode**. The Open dialog box appears.
3. In the Open dialog box, locate and select the file containing the microcode version number indicated by your release notes. (Default microcode location on Profile XP systems is *c:\profile\microcode* directory)
4. Click **Open.** The Progress Report window appears showing the microcode update task and the percentage completion.

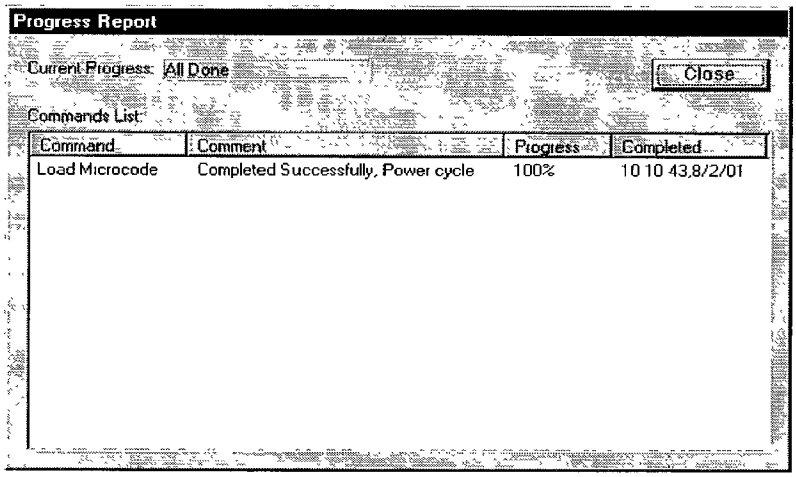

5. If you are updating a PFR500 controller, you must perform the following steps, otherwise, proceed to step 6.

   The following steps will restore default disk drive settings, which is required after updating controller firmware.

   a. Right-click the controller in the tree view, then select **Restore Disk Settings** in the context menu. The hour glass appears while the operation is performed. This takes several seconds.

b. On receiving "Operation Successful" message, click **OK**.

6. Click **File | Exit** to close the GVG Disk Utility.
7. Turn off the power to ***all attached Fibre channel RAID chassis***. Failure to shut down the RAID units will prevent the new microcode from taking effect.
8. Turn on the power to all the RAID chassis, wait three to four minutes
9. Reboot the Profile XP system.

This concludes the load controller firmware procedure.

Performing PFR500 maintenance tasks

Using the GVG Disk Utility, you can perform maintenance tasks unique to the PFR 500 storage system.

These maintenance tasks include:

Forcing a LUN online

Controllers may fail to recognize valid LUNs when expansion chassis are powered on before the controller is ready to recognize them. The LUN appears under the Bound node in the tree view, however, the LUN status is offline. Using the Disk Utility, you can force a LUN online.

To force a LUN online:

1. In the GVG Disk Utility, right-click the offline LUN, then select **Advanced | Force Online** in the context menu.
2. On receiving operation successful message, click **OK**.
3. Select the LUN in the tree view, then verify the LUN status in the right-hand status pane is is "Online".

Forcing a replacement drive to rebuild

In the event of a disk module failure, you'll repair the system by replacing the disk module as soon as possible. Sometimes, after replacing the failed drive it may be required to force the new disk module to rebuild.

To force a replacement disk module to rebuild:

1. In the GVG Disk Utility, expand the tree view to view the LUN containing the replacement drive.
2. Right-click the replacement disk, then select **Advanced | Force Rebuild** in the context menu. The rebuild process begins. (You can verify this by checking disk status in the right-hand status pane.)

   On completion, the disk module status changes from rebuild to online. You may need to refresh the GVG Disk Utility display.

Starting and Stopping disk initialization

Initializing the disk array aligns the physical disk drive heads and, where applicable, sets the striping and parity across the disk drives in a newly created logical unit. It takes a very long time to complete and is rarely needed. Initialize disk modules only when instructed to do so by Grass Valley Group support.

Restoring RAID Controller default settings

RAID Controller settings are checked every time the GVG Disk Utility is started. If for any reason, the RAID controller has inappropriate settings for operation in Profile XP installations, you must restore default RAID controller settings.

To restore RAID Controller default settings:

1. In the GVG Disk Utility tree view, right-click the RAID Controller node, then select **Advanced | Check Default Settings** in the context menu. You may receive a message asking permission to restore default settings. If so, click **OK** to continue now.
2. On receiving operation successful message, click **OK**.

   RAID Controller status is now reported as "Correct" in the right-hand status pane.

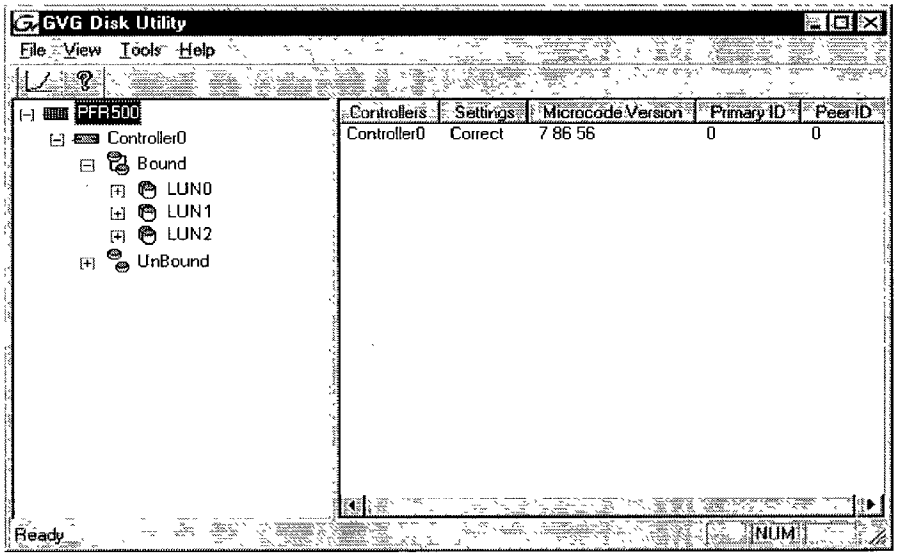

3. Reboot the PFR500 RAID Chassis containing the affected controller.
4. Reboot the Profile XP System.

Restoring disk drive default settings

If for any reason, the disk drive settings are inappropriate for operation in Profile XP installations, you must restore default disk drive settings.

1. In the GVG Disk Utility tree view, right-click the RAID Controller node, then select **Advanced | Restore Disk Settings** in the context menu.
2. On receiving an operation successful message, click **OK**.

   All disk settings on this RAID Controller are now correct.
3. Turn off the power to ***all attached Fibre channel RAID chassis***. Failure to shut down the RAID units will prevent the changes from taking effect.
4. Turn on the power to all the RAID chassis, wait three to four minutes
5. Reboot the Profile XP system.

   This concludes the restore disk settings procedure.

Reseting the PFC500 system clock

Use the following procedure to set the system clock on all PFC500 RAID Storage chassis connected. This procedure sets the PFC500 RAID controller clock to the Profile XP system clock (or the clock on which the Disk Utility is running).

To reset the PFC500 clock:

1. In the GVG Disk Utility tree view, right-click the PFC500 node, then select **Reset Clock** in the context menu.
2. On receiving an operation successful message, click **OK**.